REMARKS

Claims 19, 21-32, 34, 35, and 37 are pending. Claims 19, 21, 24, 25, 30, 32, and 35 are currently amended. No new matter is introduced by virtue of the claim amendments. Reconsideration and allowance of the present application in view of the claim amendments and remarks to follow is respectfully requested.

Telephone Conversation With Examiner

Examiner Bilgrami is thanked for the telephone conversation conducted on December 23, 2008. Proposed amendments were discussed. Cited are was discussed. No agreements were reached.

Claim Objections

Claims 19, 24, and 32 are objected to for the informalities noted on page 2 of the Office Action. In particular, claim 24 is objected to for depending on cancelled claim 20. In response, claim 24 (as well as claim 21) is amended to depend from claim 19.

Moreover, claims 19 and 32 are objected to for reciting "**by** the network access module" which the Examiner believes should recite "**at** the network access module". This objection is respectfully traversed. To begin, the basis for this objection is not clear, as the various steps of detecting, sending, receiving, caching, establishing as recited in claims 19 and 32 are actions that are performed by the network access module (NAM), not actions that occur at the NAM. In any event, for clarification, claims 19 and 32 (as well as claim 35) are amended to recite *wherein the detecting, sending, receiving, caching, and establishing are performed by the NAM to provide seamless fail-over connectivity in a manner transparent to the client application.* Accordingly, withdrawal of the claim objections is respectfully requested.

Claim Rejections – 35 USC § 112

Claims 19, 32 and 35 are rejected under 35 U.S.C. §112, first paragraph, as allegedly failing to comply with the enablement requirement. The Office Action contends that the current specification fails to disclose the claim limitations of:

(i) sending a Server Resolution Protocol request to the cluster of servers, *wherein the request includes the combination of VIA cluster name and the VIA server name* and

(ii) *cache contains mapping between the network address of the second server and the combination of the VIA cluster name and the VIA server name.*

The rejection is respectfully traversed. The specification, in fact, does disclose the above-mentioned features, for example, on page 8, line 6, through page 9, line 7, on page 10, lines 1-4. In this regard, the basis for the rejection under 112, first paragraph, is not clear. However, to the extent that the rejection is based on the specific use of the claim terms "VIA cluster name" and "VIA server name", although such terms are supported, the claims have been amended to recite "cluster name" and "server name". Accordingly, withdrawal of the rejection under 35 U.S.C. § 112, first paragraph, is respectfully requested.

Claim Rejections – 35 USC § 103

Claims 19, 21, 22, 23, 24, 25-32, 34, 35, 36, and 37 are rejected under 35 U.S.C. §103(a) as being unpatentable over Bruck et al. (U.S. 6,801,949) in view of Hart (U.S. 6,154,765) and Brendel et al. (U.S. 5,774,660). The rejections are respectfully traversed. At the very least, independent claims 19, 32 and 35 are patentable and allowable over the combination of Bruck, Hart and Brendel, for at least the following reasons.

For example, with regard to independent claims 19, 32 and 35, it is respectfully submitted that the combination of Bruck, Hart and Brendel does not disclose or suggest any of the claimed features of *detecting a failure of a first VIA connection …, sending a Server Resolution Protocol request …, receiving a Server Resolution Protocol response …, caching the response*

…, and *establishing a second VIA connection* …., wherein the *detecting, the sending, the receiving, caching, and the establishing* *are performed by the NAM to provide seamless fail-over connectivity in a manner* *transparent to the client application,* as claimed in claims 19, 32 and 35.

In formulating the rejection, the Office Action relies primarily on Bruck as disclosing the claimed features of *detecting a failure of a first VIA connection …, sending a Server Resolution Protocol request …, receiving a Server Resolution Protocol response …, caching the response …,* and *establishing a second VIA connection …..,* as recited in claims 19, 32 and 35. It is respectfully submitted that reliance on Bruck is wholly misplaced.

Bruck generally discloses a distributed server system having multiple machines that function as a front server layer between a network (such as the Internet) and a back-end server layer having multiple machines functioning as Web, FTP and/or Application servers. The front server layer is a server cluster that performs fail-over for both server layers such that when a server failure at either layer is detected, the system automatically shifts network traffic from the failed machine to one or more operational machines, reconfiguring front-layer servers as needed without interrupting operation of the server system and such that network availability is substantially unchanged regardless of machine failures (See Abstract; Col. 3, lines 19-40; and Col. 6, lines 14-60).

Bruck discloses that in contrast to conventional systems, where moving an IP address from one machine to another causes a client-server TCP connection to be lost if the machine is an end-point of the TCP connection, Bruck's method of dynamic address reassignment operates with no loss in client-server TCP connection. Bruck discloses that a server performs dynamic address reassignment using an application driver 408 (FIG. 4) to keep track of IP address movements so that old data traffic intended for an old client-server TCP connection is forwarded to the old server machine connection until the network user terminates the connection. The

dynamic address reassignment is performed using a gratuitous ARP (address resolution protocol) message, where the application driver of the server (408) sends a gratuitous ARP message packet to a router upstream (Internet side) of the server, to update the ARP cache of the appropriate routers. (See, e.g., Col. 27, lines 22-45).

In view of the above, although Bruck discloses fail over support in a clustered system, Bruck discloses that client connectivity in the case of a fail-over at a server cluster is implemented "server-side" by the server cluster through dynamic traffic network reassignment functions in which client connectivity is maintained in case of server failure without breaking network communications between clients and servers. As such, on a fundamental level, Bruck clearly does not teach or remotely suggest that *detecting, the sending, the receiving, caching, and the establishing are performed by the NAM to provide seamless fail-over connectivity in a manner transparent to the client application,* as recited in claims 19, 32 and 35. In other words, in the context of the claimed subject matter viewed as a whole, Bruck clearly does not disclose or suggest client-side fail-over support mechanism implemented by a NAM to enable seamless client connectivity to a server cluster, as claimed.

In fact, misplaced reliance on Bruck is underscored by the admission in the Office Action (page 4) that Bruck does not disclose the claimed features of establishing a VIA protocol connection or that the *detecting, sending, receiving and establishing* are performed by a NAM transparent to the client application, as claimed in claims 19, 32 and 35. Instead, to cure the deficiencies of Bruck in this regard, the Office Action cites Hart (Col. 2, lines 21-25 and Col. 8, lines 31-33) as disclosing these features. It is respectfully submitted that reliance on Hart is misplaced for various reasons.

First of all, although Hart generally discloses (in Col. 8, lines 31-33) VIA protocol communication, Hart suggests the use of VIA for communication between processing nodes in the server system, and not establishing VIA connections by a client side NAM between a client

application and a server, as claimed. Moreover, although Hart generally discloses in the Background section (Col. 2, lines 21-25) that "*[i]n a clustering system, the network client must have reconnection smarts so that the user cannot tell that behind the scenes a current connection to a server failed, and a new connection to the same IP address on another server has occurred,*" this general statement clearly does not disclose or remotely suggest, or otherwise cure the deficiencies of Bruck as noted above, with regard to *detecting, sending, receiving and establishing* being performed by a client side NAM, much less being performed by a client-side NAM to provide *seamless fail-over connectivity client transparent to the client application*, as claimed in claims 19, 32 and 35.

In fact, Hart does not even specifically disclose methods for client connectivity *per se*, much less a client side NAM providing fail over mechanisms for client connectivity to a server cluster. In contrast, Hart merely discloses a distributed rule processing environment in which an application is compiled through a specialized compiler that creates a number of digital rules to be executed in parallel over a number of processing nodes. If one of the processing nodes crashes, the distributed rule processing environment is designed to continue executing the program without the loss of data (see, Col. 4 lines 28-37).

Another fundamental flaw in the rejection of claims 19, 32 and 35, is that the Office Action (top of page 5) is devoid of any explanation providing the basis or motivation for modifying Bruck with the general, disparate teachings of Hart regarding a *network client having reconnection smarts*, to provide client-side functionality by an NAM for seamless fail over connectivity, as claimed. In any event, given the undisputed fact that Bruck teaches a server side fail over mechanism, there would appear to be no rational basis in law or fact to modify Bruck's teachings regarding "server side" fail over mechanisms with the general teachings of Hart to provide "client side" fail over functionality as claimed, as such modification would change the entire principal of operation of Bruck's fail-over system.

The impropriety of the obviousness rejections is further underscored by the fact that Bruck does not even disclose or fairly suggest the claimed features of *detecting, sending, receiving and establishing,* as recited in claims 19, 32 and 35.

For example, the Office Action contends that Bruck discloses in Col. 2, lines 38-65 *detecting a failure of a first Virtual Interface Architecture (VIA) protocol connection with the first server . . .,* as claimed. However, it is respectfully asserted that such contention is erroneous as a matter of fact. In stark contrast, Bruck discloses that in the case of server failure, the server system maintains client connections using a dynamic reconfiguration protocol that permits reassignment of network addresses in a manner that allows the server cluster to perform such operations without breaking network communications between clients and servers. (See Col. 2, lines 58-63). As such, given that client connectivity is maintained in circumstanced of server failure, the server cluster of Bruck does not detect failures of client-server connections for implementing fail-over support.

The Office Action further contends that Bruck discloses (in Col. 15, lines 20-65) *sending a Server Resolution Protocol request to the cluster of servers . . . ,* as claimed. The Examiner contends in the Response to Arguments on page 10 of the Office Action that the "Server Resolution Protocol" as disclosed on page 8, lines 26-30 and page 9, lines 1-22 of the current specification is identical to the "Address Resolution Protocol" (ARP) disclosed by Bruck. It is respectfully asserted that such contention is factually erroneous in general, as well as in the context of the claimed inventions.

Bruck discloses (in Col. 15, lines 25-39) that ARP is a conventional scheme for translating logical IP addresses into physical network interface addresses in conjunction with stored address resolution information, where network interface addresses are Media Access Control (MAC) addresses for network cards. Moreover, it is well known that ARP is a "Link Layer" protocol. In contrast, Server Resolution Protocol is an application-level protocol used for

transferring requests and responses between a client and a specific machine to ,e.g., enable the client to determine communication endpoint information of a particular server instance.

In view of the above, in the context of the claimed inventions, Bruck does not disclose or suggest *sending a Server Resolution Protocol request to the server cluster requesting connection information for a server associated with the server name and the cluster name of the server cluster,* as claimed in claims 19, 32 and 35. Indeed, as noted above, Bruck discloses that a server performs dynamic address reassignment by sending a gratuitous ARP message packet to a router that is upstream (Internet side) of the server to update the ARP cache of the appropriate routers and allow network data to be re-directed to the appropriate node. (See, e.g., Col. 27, lines 22-45). A server sending a gratuitous ARP message is not the same or similar to sending a SRP request to a server for connection information.

The Office Action also acknowledges that Bruck does not disclose caching the response from a server on a client computer, but relies on Brendel (Col. 4, lines 5-16) as disclosing the same. Brendel discloses (Col. 4, lines 5-16) that a browser can cache an IP address from a DNS server until the browser application is closed. However, Brendel discloses that the browser will not know of a server crash and can still attempt to access a crashed server after the crash has occurred using a previously cached IP address for that server. In such instance, the browser will receive no response from crashed server.

In view of the above, it is respectfully asserted that reliance on Brendel is misplaced. Although Brendel discloses caching an IP address, the cited section of Brendel does not disclose or suggest *caching a response* which, in the context of the claimed inventions providing client connectivity in fail over situations, relates to caching *a Server Resolution Protocol response from the server cluster comprising connection information of the second server*, as recited in claims 19, 32 and 35. Indeed, in the cited section of Brendel noted above, the cached IP address is not obtained in response to a request by the browser for connection information for a new working

server in the event of a crash, but merely an IP address provided by a DNS server that may be used by the browser to unknowingly access a crashed server.

Accordingly, in view of the numerous deficiencies of the cited art of record discussed above, it is respectfully submitted that claims 19, 32 and 35 are clearly patentable over the combination of Bruck, Hart and Brendel. Moreover, all pending claims depending from claims 19, 32 and 35 are patentable over Bruck, Hart and Brendel at least for the same reasons given for their respective base claims 19, 32 and 35. It is to be noted that Applicants generally deny, and do not concede to, any statement, position or averment in the Office Action in support of the claim rejections under 35 U.S.C. §103, which is not specifically addressed by the foregoing arguments and response. Withdrawal of the rejections under 35 U.S.C. § 103 is respectfully requested.

CONCLUSION

The Applicants believe that the present remarks are responsive to each of the points raised by the Examiner in the official action, and respectfully submit that all claims are in condition for allowance. Favorable consideration and passage to issue of the application at the Examiner's earliest convenience is earnestly solicited.

Date: January 15, 2009

/**Joseph F. Oriti**/
Joseph F. Oriti
Registration No. 47,835

Woodcock Washburn LLP
Cira Centre
2929 Arch Street, 12th Floor
Philadelphia, PA 19104-2891
Telephone: (215) 568-3100
Facsimile: (215) 568-3439